

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES**

In re Patent Application of)	
)	
Sanjay GHEMAWAT et al.)	ATTN: APPEAL BRIEF - PATENTS
)	
Application No.: 10/608,139)	Group Art Unit: 2163
)	
Filed: June 30, 2003)	Examiner: H. Thai
)	
For: SYSTEMS AND METHODS FOR)	
REPLICATING DATA)	

U.S. Patent and Trademark Office
Customer Window, Mail Stop Appeal Brief - Patents
Randolph Building
401 Dulany Street
Alexandria, VA 22314

APPEAL BRIEF

This Appeal Brief is submitted in response to the final Office Action, dated November 1, 2006, and in support of the Notice of Appeal, filed March 1, 2007, the period for filing the Appeal Brief being extended to June 1, 2007, by filing a petition for an extension of time and the requisite fee concurrently herewith.

I. **REAL PARTY IN INTEREST**

The real party in interest in this appeal is Google Inc.

II. **RELATED APPEALS, INTERFERENCES, AND JUDICIAL PROCEEDINGS**

Appellants are unaware of any related appeals, interferences, or judicial proceedings.

III. STATUS OF CLAIMS

Claims 1-27 are pending in this application.

Claims 1-8 and 19-27 stand rejected under 35 U.S.C. § 101 as allegedly lacking patentable utility.

Claims 1-8 stand rejected under 35 U.S.C. § 102(b) as allegedly anticipated by Jindal et al. (U.S. Patent No. 6,324,580).

Claims 19-27 stand rejected under 35 U.S.C. § 103(a) as allegedly unpatentable over Jindal et al. and Narendran et al. (U.S. Patent No. 6,070,191).

Claims 9-18 stand withdrawn from consideration as a result of an election to a restriction requirement.

Claims 1-8 and 19-27 are the subject of the present appeal. These claims are reproduced in the Claim Appendix of this Appeal Brief.

IV. STATUS OF AMENDMENTS

An After Final Request for Reconsideration was filed on December 29, 2006, subsequent to the final Office Action. No claim amendments were filed subsequent to the final Office Action. The Examiner issued an Advisory Action, dated January 30, 2007, that indicated that the After Final Request for Reconsideration did not place the application in condition for allowance and would not be entered for purposes of appeal. A Pre-Appeal Brief Request for Review was filed on March 1, 2007, which resulted in a finding that there is allegedly at least one actual issue for appeal.

V. SUMMARY OF CLAIMED SUBJECT MATTER

In the paragraphs that follow, a concise explanation of the independent claims and the claims reciting means-plus-function or step-plus-function language that are involved in this appeal will be provided by referring, in parenthesis, to examples of where support can be found in the specification and drawings.

Claim 1 recites a method for distributing data in a system (Fig. 1, FILE SYSTEM) that includes a plurality of servers (Fig. 1, 120). The method comprises identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers (Fig. 6, 620; paragraph 0063), prior data distribution involving the servers (Fig. 6, 630; paragraph 0064), or failure correlation properties associated with the servers (Fig. 6, 640; paragraph 0065); and placing the replicas of the data at the identified servers (Fig. 6, 650; paragraph 0066).

Claim 2 recites that identifying ones of the servers includes identifying underutilized ones of the servers as candidates to store the replicas of the data (Fig. 6, 620; paragraph 0063).

Claim 4 recites that identifying ones of the servers includes identifying ones of the servers that have not been involved in a recent data distribution as candidates to store the replicas of the data (Fig. 6, 630; paragraph 0064).

Claim 5 recites that identifying ones of the servers includes identifying system conditions that affect two or more of the servers, and identifying ones of the servers as candidates to store the replicas of the data based on the identified system conditions (Fig. 6, 640; paragraph 0065).

Claim 6 recites that the number of the replicas of the data stored by the servers is user-configurable (paragraphs 0034 and 0068).

Claim 7 recites a system (Fig. 1, FILE SYSTEM) for distributing chunks in a network (Fig. 1, 100) that includes a plurality of servers (Fig. 1, 120). The system comprises means for selecting (Fig. 1, 130) ones of the servers to store replicas of the chunks based on at least one of utilization of the servers (Fig. 6, 620; paragraph 0063), prior chunk distribution involving the servers (Fig. 6, 630; paragraph 0064), or failure correlation properties associated with the servers (Fig. 6, 640; paragraph 0065); and means for storing (Fig. 1, 130) the replicas of the chunks at the selected servers (Fig. 6, 650; paragraph 0066) to at least one of increase reliability of the chunks (paragraph 0067), increase availability of the chunks (paragraph 0067), or increase bandwidth utilization in the system (paragraph 0067).

Claim 8 recites a file system (Fig. 1, FILE SYSTEM) comprises a plurality of servers (Fig. 1, 120) that store replicas of chunks (paragraph 0034); and a master (Fig. 1, 130) connected to the servers. The master is configured to identify one or more of the servers to store a replica of a chunk based on at least one of utilization of the servers (Fig. 6, 620; paragraph 0063), prior chunk distribution involving the servers (Fig. 6, 630; paragraph 0064), or failure correlation properties associated with the servers (Fig. 6, 640; paragraph 0065), and place the replicas of the chunk at the identified one or more servers (Fig. 6, 650; paragraph 0066).

Claim 19 recites a method for redistributing chunks of data in a system (Fig. 1, FILE SYSTEM) that includes a plurality of servers (Fig. 1, 120) that store replicas of the chunks. The method comprises monitoring utilization of the servers (Fig. 8, 810; paragraph 0072); determining whether to redistribute any of the replicas (Fig. 8, 820; paragraph 0072); selecting one or more of the replicas to redistribute based on the utilization of the servers (Fig. 8, 820; paragraph 0072); selecting one or more of the servers to which to move the one or more replicas

(Fig. 8, 830-850; paragraph 0073); and moving the one or more replicas to the selected one or more servers (paragraph 0073).

Claim 21 recites identifying underutilized ones of the servers as candidates to which to move the one or more replicas (paragraphs 0063 and 0073).

Claim 23 recites identifying ones of the servers that have not been involved in a recent redistribution as candidates to which to move the one or more replicas (paragraphs 0064 and 0073).

Claim 24 recites determining failure correlation properties associated with the servers (paragraphs 0065 and 0073), and identifying ones of the servers based on the failure correlation properties as candidates to which to move the one or more replicas (paragraphs 0065 and 0073).

Claim 25 recites deleting the one or more replicas from one or more of the servers (paragraph 0073), and instructing the selected one or more servers to copy the one or more replicas from another one or more of the servers (paragraph 0073).

Claim 26 recites a system (Fig. 1, FILE SYSTEM) for redistributing data in a network (Fig. 1, 100) that includes a plurality of servers (Fig. 1, 120) that store replicas of the data. The system comprises means for monitoring (Fig. 1, 130) utilization of the servers (Fig. 8, 810; paragraph 0072); means for selecting (Fig. 1, 130) one or more of the replicas to redistribute based on the utilization of the servers (Fig. 8, 820; paragraph 0072); means for identifying (Fig. 1, 130) one or more of the servers to which to move the one or more replicas (Fig. 8, 830-850; paragraph 0073); and means for redistributing (Fig. 1, 130) the one or more replicas to the identified one or more servers (paragraph 0073) to at least one of increase reliability of the data (paragraphs 0067 and 0073), increase availability of the data (paragraphs 0067 and 0073), or

increase bandwidth utilization in the system (paragraphs 0067 and 0073).

Claim 27 recites a file system (Fig. 1, FILE SYSTEM) comprises a plurality of servers (Fig. 1, 120) configured to store replicas of chunks of data (paragraph 0034); and a master (Fig. 1, 130) connected to the servers. The master is configured to select one or more of the replicas to redistribute based on utilization of the servers (Fig. 8, 820; paragraph 0072), identify one or more of the servers to which to move the selected one or more replicas (Fig. 8, 830-850; paragraph 0073), and move the selected one or more replicas to the identified one or more servers (paragraph 0073).

VI. GROUND OF REJECTION TO BE REVIEWED ON APPEAL

- A. Claims 1-8 and 19-27 stand rejected under 35 U.S.C. § 101 as lacking patentable utility.
- B. Claims 1-8 stand rejected under 35 U.S.C. § 102(b) as anticipated by Jindal et al.
- C. Claims 19-27 stand rejected under 35 U.S.C. § 103(a) as unpatentable over Jindal et al. and Narendran et al.

VII. ARGUMENT

A. The Rejection Under 35 U.S.C. § 101 Should be Reversed.

The initial burden of establishing a prima facie basis to deny patentability to a claimed invention is always upon the Examiner. In re Oetiker, 977 F.2d 1443, 24 USPQ2d 1443 (Fed. Cir. 1992).

- 1. Claims 1-6.

Independent claim 1 is directed to a method for distributing data in a system that includes a plurality of servers. The method comprises identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers; and placing the replicas of the data at the identified servers.

The Examiner rejected claim 1 under 35 U.S.C. § 101 as allegedly lacking patentable utility. Final Office Action, paragraph 5. The Examiner alleged that claim 1 merely identifies or defines "a data process in which to be manipulated without giving rise to a concrete, useful and tangible result." Final Office Action, page 3. Appellants traverse the rejection and submit that the Examiner's allegation falls short of establishing a *prima facie* case to deny patentability under 35 U.S.C. § 101.

The USPTO specifically states, in the "Interim Guidelines for Examination of Patent Applications for Patent Subject Matter Eligibility" (hereinafter "Guidelines"), the following:

To properly determine whether a claimed invention complies with the statutory invention requirements of 35 U.S.C. 101, USPTO personnel must first identify whether the claim falls within at least one of the four enumerated categories of patentable subject matter recited in section 101 (process, machine, manufacture or composition of matter).

* * *

The burden is on the USPTO to set forth a prima facie case of unpatentability. Therefore if the examiner determines that it is more likely than not that the claimed subject matter falls outside all of the statutory categories, the examiner must provide an explanation. ... If the examiner can establish a prima facie case that a claim does not fall into a statutory category, that does not preclude complete examination of the application for satisfaction of all other conditions of patentability. The examiner must further continue with the statutory subject matter analysis as set forth below.

Interim Guidelines For Examination of Patent Applications for Patent Subject Matter Eligibility, 1300 OG 142, § IV. B., November 22, 2005. Despite this clear obligation, the Examiner failed to make a determination of whether claim 1 falls within one of the four enumerated categories of

patentable subject matter recited in Section 101. Appellants respectfully submit that claim 1 falls within one of the four enumerated categories of patentable subject matter recited in Section 101.

For example, claim 1 is directed to a method for distributing data in a system that includes a plurality of servers. Thus, claim 1 is directed to a "process," which is one of the four enumerated categories of patentable subject matter recited in Section 101. The Examiner's allegation that the claim 1 identifies or defines a data process falls short of establishing that the claim falls outside all of the statutory categories, because a "process" is one of the statutory categories.

In light of the foregoing reasons, Appellants submit that the rejection of claim 1 under 35 U.S.C. § 101 is improper.

Furthermore, even assuming, arguendo, that claim 1 is not directed to one of the four enumerated categories of patentable subject matter recited in 35 U.S.C. § 101, the Guidelines further state:

While abstract ideas, natural phenomena, and laws of nature are not eligible for patenting, methods and products employing abstract ideas, natural phenomena, and laws of nature to perform a real-world function may well be. In evaluating whether a claim meets the requirements of section 101, the claim must be considered as a whole to determine whether it is for a particular application of an abstract idea, natural phenomenon, or law of nature, rather than for the abstract idea, natural phenomenon, or law of nature itself.

Id. at § IV. C.

A proper analysis under 35 U.S.C. § 101 entails a determination of whether the claimed invention is a practical application of an abstract idea, law of nature, or natural phenomenon (judicial exceptions to 35 U.S.C. § 101). The Examiner has not even alleged that Appellants' claim 1 is directed to an abstract idea, natural phenomenon, or a law of nature. Therefore, the

Examiner's allegation that claim 1 identifies or defines a data process falls short of establishing that the claim is directed to an abstract idea, a natural phenomenon, or a law of nature.

Nevertheless, claim 1 recites a practical application because it provides a transformation of an article to a different state or thing. The Guidelines state that if a claim provides a transformation or reduction of an article to a different state or thing, then the Examiner shall end the inquiry and find that the claim meets the statutory requirement of Section 101. Id. at IV.C.2.a. Claim 1 recites placing replicas of data at the identified servers, thus transforming a collection of data into replicas of data that are placed (i.e., stored) at identified servers.

In light of this, Appellants submit that the rejection of claim 1 under 35 U.S.C. § 101 is again improper.

Even assuming, *arguendo*, that claim 1 does not provide a transformation, claim 1 provides a practical application that produces a useful, tangible, and concrete result. According to the Guidelines, in determining whether a claim provides a practical application that produces a useful, tangible, and concrete result, the Examiner should consider and weigh the following factors: (1) the USPTO's official interpretation of the utility requirement provides that the utility of an invention has to be (i) specific, (ii) substantial and (iii) credible; (2) the tangible requirement does not necessarily mean that a claim must either be tied to a particular machine or apparatus or must operate to change articles or materials to a different state or thing; and (3) the process must have a result that can be substantially repeatable or the process must substantially produce the same result again. Id. at IV.C.2.b.

The Examiner did not consider any of these factors. Instead, the Examiner alleged that Appellants' claim 1 does not "do an action statement" because Appellants' claim language of

"identifying servers and placing the replicas of the data at the identified servers does not render a result." Final Office Action, page 2. Appellants submit that the Examiner's statement lacks merit.

Claim 1 recites "placing the replicas of the data at the identified servers." The result clearly is that the replicas of the data are placed (i.e., stored) at the identified servers. This result is clearly "useful" in the sense that storage of the replicas at the identified servers maximizes data reliability and availability, and maximizes network utilization. See, e.g., Appellants' specification at paragraph 0067. This result is also clearly "tangible" in the sense that particular replicas are placed (i.e., stored) at particular servers. See, e.g., Appellants' specification at paragraph 0067. This result is clearly "concrete" in the sense that data is placed at the identified servers and this placement can be repeated. See, e.g., Appellants' specification at paragraphs 0063-0066. Accordingly, the Examiner's statement that claim 1 lacks a result is without merit.

Further, the Examiner indicated that Appellants "need to use language such as in dependent claim 4: use of 'identifying servers to store replicas of the data.'" Final Office Action, page 2. Appellants note that language very similar to the language of "identifying servers to store replicas of the data" is already recited in claim 1. Therefore, claim 1, under the Examiner's reasoning, should satisfy the requirements of 35 U.S.C. § 101. It appears that under the guise of a rejection under 35 U.S.C. § 101, the Examiner is attempting to cure deficiencies in the Examiner's rejections under 35 U.S.C. § 102 and 103 by trying to force Appellants to amend the claims.

For at least these reasons, it is respectfully submitted that claims 1-6 recite statutory subject matter under 35 U.S.C. § 101. Reversal of the rejection of claims 1-6 is respectfully

requested.

2. Claim 7.

Independent claim 7 is directed to a system for distributing chunks in a network that includes a plurality of servers. The system comprises means for selecting ones of the servers to store replicas of the chunks based on at least one of utilization of the servers, prior chunk distribution involving the servers, or failure correlation properties associated with the servers; and means for storing the replicas of the chunks at the selected servers to at least one of increase reliability of the chunks, increase availability of the chunks, or increase bandwidth utilization in the system.

The Examiner rejected claim 7 under 35 U.S.C. § 101 as allegedly lacking patentable utility. Final Office Action, paragraph 5. The Examiner provided no evidence to support the broad allegation that claim 7 lacks patentable utility. Id. Thus, the Examiner did not establish a prima facie case to support the broad allegation that claim 7 lacks patentable utility.

Accordingly, Appellants traverse the rejection and submit that the Examiner's allegation falls short of establishing a prima facie case to deny patentability under 35 U.S.C. § 101.

The USPTO specifically states, in the "Interim Guidelines for Examination of Patent Applications for Patent Subject Matter Eligibility" (hereinafter "Guidelines"), the following:

To properly determine whether a claimed invention complies with the statutory invention requirements of 35 U.S.C. 101, USPTO personnel must first identify whether the claim falls within at least one of the four enumerated categories of patentable subject matter recited in section 101 (process, machine, manufacture or composition of matter).

* * *

The burden is on the USPTO to set forth a prima facie case of unpatentability. Therefore if the examiner determines that it is more likely than not that the claimed subject matter falls outside all of the statutory categories, the examiner must provide an explanation. ... If the examiner can establish a prima facie case that a claim does not fall into a statutory

category, that does not preclude complete examination of the application for satisfaction of all other conditions of patentability. The examiner must further continue with the statutory subject matter analysis as set forth below.

Interim Guidelines For Examination of Patent Applications for Patent Subject Matter Eligibility, 1300 OG 142, § IV. B., November 22, 2005. Despite this clear obligation, the Examiner failed to make a determination of whether claim 7 falls within one of the four enumerated categories of patentable subject matter recited in Section 101. Appellants respectfully submit that claim 7 falls within one of the four enumerated categories of patentable subject matter recited in Section 101.

For example, claim 7 is directed to a system for distributing chunks in a network that includes a plurality of servers. Thus, claim 7 is directed to a "machine," which is one of the four enumerated categories of patentable subject matter recited in Section 101.

In light of the foregoing reasons, Appellants submit that the rejection of claim 7 under 35 U.S.C. § 101 is improper.

Furthermore, even assuming, arguendo, that claim 7 is not directed to one of the four enumerated categories of patentable subject matter recited in 35 U.S.C. § 101, the Guidelines further state:

While abstract ideas, natural phenomena, and laws of nature are not eligible for patenting, methods and products employing abstract ideas, natural phenomena, and laws of nature to perform a real-world function may well be. In evaluating whether a claim meets the requirements of section 101, the claim must be considered as a whole to determine whether it is for a particular application of an abstract idea, natural phenomenon, or law of nature, rather than for the abstract idea, natural phenomenon, or law of nature itself.

Id. at § IV. C.

A proper analysis under 35 U.S.C. § 101 entails a determination of whether the claimed invention is a practical application of an abstract idea, law of nature, or natural phenomenon

(judicial exceptions to 35 U.S.C. § 101). The Examiner has not even alleged that Appellants' claim 7 is directed to an abstract idea, natural phenomenon, or a law of nature. Therefore, the Examiner's allegation that claim 7 identifies or defines a data process falls short of establishing that the claim is directed to an abstract idea, a natural phenomenon, or a law of nature.

Nevertheless, claim 7 recites a practical application because it provides a transformation of an article to a different state or thing. The Guidelines state that if a claim provides a transformation or reduction of an article to a different state or thing, then the Examiner shall end the inquiry and find that the claim meets the statutory requirement of Section 101. Id. at IV.C.2.a. Claim 7 recites means for storing the replicas of the chunks at the selected servers, thus transforming a collection of data into replicas of data that are stored by selected servers.

In light of this, Appellants submit that the rejection of claim 7 under 35 U.S.C. § 101 is again improper.

Even assuming, arguendo, that claim 7 does not provide a transformation, claim 7 provides a practical application that produces a useful, tangible, and concrete result. According to the Guidelines, in determining whether a claim provides a practical application that produces a useful, tangible, and concrete result, the Examiner should consider and weigh the following factors: (1) the USPTO's official interpretation of the utility requirement provides that the utility of an invention has to be (i) specific, (ii) substantial and (iii) credible; (2) the tangible requirement does not necessarily mean that a claim must either be tied to a particular machine or apparatus or must operate to change articles or materials to a different state or thing; and (3) the process must have a result that can be substantially repeatable or the process must substantially produce the same result again. Id. at IV.C.2.b.

Claim 7 recites "means for storing the replicas of the chunks at the selected servers to at least one of increase reliability of the chunks, increase availability of the chunks, or increase bandwidth utilization in the system." The result clearly is that the replicas of the data are stored at the selected servers. This result is clearly "useful" in the sense that storage of the replicas at the selected servers increases reliability of the chunks, increases availability of the chunks, and/or increases bandwidth utilization in the system. This result is also clearly "tangible" in the sense that particular replicas are stored at particular servers. See, e.g., Appellants' specification at paragraph 0067. This result is clearly "concrete" in the sense that data is stored at the selected servers and this storage can be repeated. See, e.g., Appellants' specification at paragraphs 0063-0066. Accordingly, claim 7 clearly recites a useful, tangible, and concrete result.

For at least these reasons, it is respectfully submitted that claim 7 recites statutory subject matter under 35 U.S.C. § 101. Reversal of the rejection of claim 7 is respectfully requested.

3. Claim 8.

Independent claim 8 is directed to a file system that comprises a plurality of servers that store replicas of chunks, and a master connected to the servers. The master is configured to identify one or more of the servers to store a replica of a chunk based on at least one of utilization of the servers, prior chunk distribution involving the servers, or failure correlation properties associated with the servers, and place the replicas of the chunk at the identified one or more servers.

The Examiner rejected claim 8 under 35 U.S.C. § 101 as allegedly lacking patentable utility. Final Office Action, paragraph 5. The Examiner alleged that claim 8 merely identifies or defines "a data process in which to be manipulated without giving rise to a concrete, useful and

tangible result." Final Office Action, page 3. Appellants traverse the rejection and submit that the Examiner's allegation falls short of establishing a *prima facie* case to deny patentability under 35 U.S.C. § 101.

The USPTO specifically states, in the "Interim Guidelines for Examination of Patent Applications for Patent Subject Matter Eligibility" (hereinafter "Guidelines"), the following:

To properly determine whether a claimed invention complies with the statutory invention requirements of 35 U.S.C. 101, USPTO personnel must first identify whether the claim falls within at least one of the four enumerated categories of patentable subject matter recited in section 101 (process, machine, manufacture or composition of matter).

* * *

The burden is on the USPTO to set forth a *prima facie* case of unpatentability. Therefore if the examiner determines that it is more likely than not that the claimed subject matter falls outside all of the statutory categories, the examiner must provide an explanation. ... If the examiner can establish a *prima facie* case that a claim does not fall into a statutory category, that does not preclude complete examination of the application for satisfaction of all other conditions of patentability. The examiner must further continue with the statutory subject matter analysis as set forth below.

Interim Guidelines For Examination of Patent Applications for Patent Subject Matter Eligibility, 1300 OG 142, § IV. B., November 22, 2005. Despite this clear obligation, the Examiner failed to make a determination of whether claim 8 falls within one of the four enumerated categories of patentable subject matter recited in Section 101. Appellants respectfully submit that claim 8 falls within one of the four enumerated categories of patentable subject matter recited in Section 101.

For example, claim 8 is directed to a file system that includes a plurality of servers and a master. Thus, claim 8 is directed to a "machine," which is one of the four enumerated categories of patentable subject matter recited in Section 101. The Examiner's allegation that the claim 8 identifies or defines a data process lacks merit because claim 8 is directed to a file system and not a data process.

In light of the foregoing reasons, Appellants submit that the rejection of claim 8 under 35 U.S.C. § 101 is improper.

Furthermore, even assuming, arguendo, that claim 8 is not directed to one of the four enumerated categories of patentable subject matter recited in 35 U.S.C. § 101, the Guidelines further state:

While abstract ideas, natural phenomena, and laws of nature are not eligible for patenting, methods and products employing abstract ideas, natural phenomena, and laws of nature to perform a real-world function may well be. In evaluating whether a claim meets the requirements of section 101, the claim must be considered as a whole to determine whether it is for a particular application of an abstract idea, natural phenomenon, or law of nature, rather than for the abstract idea, natural phenomenon, or law of nature itself.

Id. at § IV. C.

A proper analysis under 35 U.S.C. § 101 entails a determination of whether the claimed invention is a practical application of an abstract idea, law of nature, or natural phenomenon (judicial exceptions to 35 U.S.C. § 101). The Examiner has not even alleged that Appellants' claim 8 is directed to an abstract idea, natural phenomenon, or a law of nature. Therefore, the Examiner's allegation that claim 8 identifies or defines a data process falls short of establishing that the claim is directed to an abstract idea, a natural phenomenon, or a law of nature.

Nevertheless, claim 8 recites a practical application because it provides a transformation of an article to a different state or thing. The Guidelines state that if a claim provides a transformation or reduction of an article to a different state or thing, then the Examiner shall end the inquiry and find that the claim meets the statutory requirement of Section 101. Id. at IV.C.2.a. Claim 8 recites a master that, among other things, places (i.e., stores) the replicas of the chunk at the identified one or more servers, thus transforming a collection of data into replicas of data that are stored by identified one or more servers.

In light of this, Appellants submit that the rejection of claim 8 under 35 U.S.C. § 101 is again improper.

Even assuming, *arguendo*, that claim 8 does not provide a transformation, claim 8 provides a practical application that produces a useful, tangible, and concrete result. According to the Guidelines, in determining whether a claim provides a practical application that produces a useful, tangible, and concrete result, the Examiner should consider and weigh the following factors: (1) the USPTO's official interpretation of the utility requirement provides that the utility of an invention has to be (i) specific, (ii) substantial and (iii) credible; (2) the tangible requirement does not necessarily mean that a claim must either be tied to a particular machine or apparatus or must operate to change articles or materials to a different state or thing; and (3) the process must have a result that can be substantially repeatable or the process must substantially produce the same result again. *Id.* at IV.C.2.b.

The Examiner did not consider any of these factors. Instead, the Examiner alleged that Appellants' claim 8 does not "do an action statement" because Appellants' claim language of "identifying servers and placing the replicas of the data at the identified servers does not render a result." Final Office Action, page 2. Appellants submit that the Examiner's statement lacks merit.

Claim 8 recites a master that is configured to "place the replicas of the chunk at the identified one or more servers." The result clearly is that the replicas of the data are placed (i.e., stored) at the identified one or more servers. This result is clearly "useful" in the sense that storage of the replicas at the identified servers maximizes data reliability and availability, and maximizes network utilization. See, e.g., Appellants' specification at paragraph 0067. This

result is also clearly "tangible" in the sense that particular replicas are stored at particular servers. See, e.g., Appellants' specification at paragraph 0067. This result is clearly "concrete" in the sense that data is placed at the identified one or more servers and this placement can be repeated. See, e.g., Appellants' specification at paragraphs 0063-0066. Accordingly, the Examiner's statement that claim 8 lacks a result is without merit.

Further, the Examiner indicated that Appellants "need to use language such as in dependent claim 4: use of 'identifying servers to store replicas of the data.'" Final Office Action, page 2. Appellants note that the language of "identifying servers to store replicas of the data" is already recited in claim 8. Therefore, claim 8, under the Examiner's reasoning, should satisfy the requirements of 35 U.S.C. § 101. It appears that under the guise of a rejection under 35 U.S.C. § 101, the Examiner is attempting to cure deficiencies in the Examiner's rejections under 35 U.S.C. § 102 and 103 by trying to force Appellants to amend the claims.

For at least these reasons, it is respectfully submitted that claim 8 recites statutory subject matter under 35 U.S.C. § 101. Reversal of the rejection of claim 8 is respectfully requested.

4. Claims 19-25.

Independent claim 19 is directed to a method for redistributing chunks of data in a system that includes a plurality of servers that store replicas of the chunks. The method comprises monitoring utilization of the servers; determining whether to redistribute any of the replicas; selecting one or more of the replicas to redistribute based on the utilization of the servers; selecting one or more of the servers to which to move the one or more replicas; and moving the one or more replicas to the selected one or more servers.

The Examiner rejected claim 19 under 35 U.S.C. § 101 as allegedly lacking patentable

utility. Final Office Action, paragraph 5. The Examiner alleged that claim 19 merely identifies or defines "a data process in which to be manipulated without giving rise to a concrete, useful and tangible result." Final Office Action, page 3. Appellants traverse the rejection and submit that the Examiner's allegation falls short of establishing a *prima facie* case to deny patentability under 35 U.S.C. § 101.

The USPTO specifically states, in the "Interim Guidelines for Examination of Patent Applications for Patent Subject Matter Eligibility" (hereinafter "Guidelines"), the following:

To properly determine whether a claimed invention complies with the statutory invention requirements of 35 U.S.C. 101, USPTO personnel must first identify whether the claim falls within at least one of the four enumerated categories of patentable subject matter recited in section 101 (process, machine, manufacture or composition of matter).

* * *

The burden is on the USPTO to set forth a prima facie case of unpatentability. Therefore if the examiner determines that it is more likely than not that the claimed subject matter falls outside all of the statutory categories, the examiner must provide an explanation. ... If the examiner can establish a prima facie case that a claim does not fall into a statutory category, that does not preclude complete examination of the application for satisfaction of all other conditions of patentability. The examiner must further continue with the statutory subject matter analysis as set forth below.

Interim Guidelines For Examination of Patent Applications for Patent Subject Matter Eligibility, 1300 OG 142, § IV. B., November 22, 2005. Despite this clear obligation, the Examiner failed to make a determination of whether claim 19 falls within one of the four enumerated categories of patentable subject matter recited in Section 101. Appellants respectfully submit that claim 19 falls within one of the four enumerated categories of patentable subject matter recited in Section 101.

For example, claim 19 is directed to a method for redistributing chunks of data in a system that includes a plurality of servers that store replicas of the chunks. Thus, claim 19 is directed to a "process," which is one of the four enumerated categories of patentable subject

matter recited in Section 101. The Examiner's allegation that the claim 19 identifies or defines a data process falls short of establishing that the claim falls outside all of the statutory categories, because a "process" is one of the statutory categories.

In light of the foregoing reasons, Appellants submit that the rejection of claim 19 under 35 U.S.C. § 101 is improper.

Furthermore, even assuming, arguendo, that claim 19 is not directed to one of the four enumerated categories of patentable subject matter recited in 35 U.S.C. § 101, the Guidelines further state:

While abstract ideas, natural phenomena, and laws of nature are not eligible for patenting, methods and products employing abstract ideas, natural phenomena, and laws of nature to perform a real-world function may well be. In evaluating whether a claim meets the requirements of section 101, the claim must be considered as a whole to determine whether it is for a particular application of an abstract idea, natural phenomenon, or law of nature, rather than for the abstract idea, natural phenomenon, or law of nature itself.

Id. at § IV. C.

A proper analysis under 35 U.S.C. § 101 entails a determination of whether the claimed invention is a practical application of an abstract idea, law of nature, or natural phenomenon (judicial exceptions to 35 U.S.C. § 101). The Examiner has not even alleged that Appellants' claim 19 is directed to an abstract idea, natural phenomenon, or a law of nature. Therefore, the Examiner's allegation that claim 19 identifies or defines a data process falls short of establishing that the claim is directed to an abstract idea, a natural phenomenon, or a law of nature.

Nevertheless, claim 19 recites a practical application because it provides a transformation of an article to a different state or thing. The Guidelines state that if a claim provides a transformation or reduction of an article to a different state or thing, then the Examiner shall end the inquiry and find that the claim meets the statutory requirement of Section 101. Id. at

IV.C.2.a. Claim 19 recites moving the one or more replicas to the selected one or more servers, thus transforming a collection of data into replicas of data that are stored by selected servers.

In light of this, Appellants submit that the rejection of claim 19 under 35 U.S.C. § 101 is again improper.

Even assuming, arguendo, that claim 19 does not provide a transformation, claim 19 provides a practical application that produces a useful, tangible, and concrete result. According to the Guidelines, in determining whether a claim provides a practical application that produces a useful, tangible, and concrete result, the Examiner should consider and weigh the following factors: (1) the USPTO's official interpretation of the utility requirement provides that the utility of an invention has to be (i) specific, (ii) substantial and (iii) credible; (2) the tangible requirement does not necessarily mean that a claim must either be tied to a particular machine or apparatus or must operate to change articles or materials to a different state or thing; and (3) the process must have a result that can be substantially repeatable or the process must substantially produce the same result again. Id. at IV.C.2.b.

The Examiner did not consider any of these factors. Instead, the Examiner alleged that Appellants' claim 19 does not "do an action statement" because Appellants' claim language of "identifying servers and placing the replicas of the data at the identified servers does not render a result." Final Office Action, page 2. Appellants submit that the Examiner's statement lacks merit.

Claim 19 recites "moving the one or more replicas to the selected one or more servers." The result clearly is that the replicas of the data are moved (and stored) to the selected servers. This result is clearly "useful" in the sense that storage of the replicas at the selected servers

maximizes data reliability and availability, and maximizes network utilization. See, e.g., Appellants' specification at paragraphs 0067 and 0073. This result is also clearly "tangible" in the sense that particular replicas are stored at particular servers. See, e.g., Appellants' specification at paragraph 0073. This result is clearly "concrete" in the sense that data is moved to the selected servers and this movement can be repeated. See, e.g., Appellants' specification at paragraphs 0072-0073. Accordingly, the Examiner's statement that claim 19 lacks a result is without merit.

Further, the Examiner indicated that Appellants "need to use language such as in dependent claim 4: use of 'identifying servers to store replicas of the data.'" Final Office Action, page 2. Appellants note that the language of "identifying servers to store replicas of the data" is similar to the language of "selecting one or more of the servers to which to move the one or more replicas" that is already recited in claim 19. Therefore, claim 19, under the Examiner's reasoning, should satisfy the requirements of 35 U.S.C. § 101. It appears that under the guise of a rejection under 35 U.S.C. § 101, the Examiner is attempting to cure deficiencies in the Examiner's rejections under 35 U.S.C. § 102 and 103 by trying to force Appellants to amend the claims.

For at least these reasons, it is respectfully submitted that claims 19-25 recite statutory subject matter under 35 U.S.C. § 101. Reversal of the rejection of claims 19-25 is respectfully requested.

5. Claim 26.

Independent claim 26 is directed to a system for redistributing data in a network that includes a plurality of servers that store replicas of the data. The system comprises means for monitoring utilization of the servers; means for selecting one or more of the replicas to redistribute based on the utilization of the servers; means for identifying one or more of the servers to which to move the one or more replicas; and means for redistributing the one or more replicas to the identified one or more servers to at least one of increase reliability of the data, increase availability of the data, or increase bandwidth utilization in the system.

The Examiner rejected claim 26 under 35 U.S.C. § 101 as allegedly lacking patentable utility. Final Office Action, paragraph 5. The Examiner alleged that claim 26 merely identifies or defines "a data process in which to be manipulated without giving rise to a concrete, useful and tangible result." Final Office Action, page 3. Appellants traverse the rejection and submit that the Examiner's allegation falls short of establishing a *prima facie* case to deny patentability under 35 U.S.C. § 101.

The USPTO specifically states, in the "Interim Guidelines for Examination of Patent Applications for Patent Subject Matter Eligibility" (hereinafter "Guidelines"), the following:

To properly determine whether a claimed invention complies with the statutory invention requirements of 35 U.S.C. 101, USPTO personnel must first identify whether the claim falls within at least one of the four enumerated categories of patentable subject matter recited in section 101 (process, machine, manufacture or composition of matter).

* * *

The burden is on the USPTO to set forth a prima facie case of unpatentability. Therefore if the examiner determines that it is more likely than not that the claimed subject matter falls outside all of the statutory categories, the examiner must provide an explanation. ... If the examiner can establish a prima facie case that a claim does not fall into a statutory category, that does not preclude complete examination of the application for satisfaction of all other conditions of patentability. The examiner must further continue with the statutory subject matter analysis as set forth below.

Interim Guidelines For Examination of Patent Applications for Patent Subject Matter Eligibility,

1300 OG 142, § IV. B., November 22, 2005. Despite this clear obligation, the Examiner failed to make a determination of whether claim 26 falls within one of the four enumerated categories of patentable subject matter recited in Section 101. Appellants respectfully submit that claim 26 falls within one of the four enumerated categories of patentable subject matter recited in Section 101.

For example, claim 26 is directed to a system for redistributing data in a network that includes a plurality of servers that store replicas of the data. Thus, claim 26 is directed to a "machine," which is one of the four enumerated categories of patentable subject matter recited in Section 101. The Examiner's allegation that the claim 26 identifies or defines a data process lacks merit because claim 26 is directed to a system and not a data process.

In light of the foregoing reasons, Appellants submit that the rejection of claim 26 under 35 U.S.C. § 101 is improper.

Furthermore, even assuming, *arguendo*, that claim 26 is not directed to one of the four enumerated categories of patentable subject matter recited in 35 U.S.C. § 101, the Guidelines further state:

While abstract ideas, natural phenomena, and laws of nature are not eligible for patenting, methods and products employing abstract ideas, natural phenomena, and laws of nature to perform a real-world function may well be. In evaluating whether a claim meets the requirements of section 101, the claim must be considered as a whole to determine whether it is for a particular application of an abstract idea, natural phenomenon, or law of nature, rather than for the abstract idea, natural phenomenon, or law of nature itself.

Id. at § IV. C.

A proper analysis under 35 U.S.C. § 101 entails a determination of whether the claimed invention is a practical application of an abstract idea, law of nature, or natural phenomenon (judicial exceptions to 35 U.S.C. § 101). The Examiner has not even alleged that Appellants'

claim 26 is directed to an abstract idea, natural phenomenon, or a law of nature. Therefore, the Examiner's allegation that claim 26 identifies or defines a data process falls short of establishing that the claim is directed to an abstract idea, a natural phenomenon, or a law of nature.

Nevertheless, claim 26 recites a practical application because it provides a transformation of an article to a different state or thing. The Guidelines state that if a claim provides a transformation or reduction of an article to a different state or thing, then the Examiner shall end the inquiry and find that the claim meets the statutory requirement of Section 101. Id. at IV.C.2.a. Claim 26 recites means for redistributing the one or more replicas to the identified one or more servers, thus transforming a collection of data into replicas of data that are redistributed and stored by one or more servers.

In light of this, Appellants submit that the rejection of claim 26 under 35 U.S.C. § 101 is again improper.

Even assuming, arguendo, that claim 26 does not provide a transformation, claim 26 provides a practical application that produces a useful, tangible, and concrete result. According to the Guidelines, in determining whether a claim provides a practical application that produces a useful, tangible, and concrete result, the Examiner should consider and weigh the following factors: (1) the USPTO's official interpretation of the utility requirement provides that the utility of an invention has to be (i) specific, (ii) substantial and (iii) credible; (2) the tangible requirement does not necessarily mean that a claim must either be tied to a particular machine or apparatus or must operate to change articles or materials to a different state or thing; and (3) the process must have a result that can be substantially repeatable or the process must substantially produce the same result again. Id. at IV.C.2.b.

The Examiner did not consider any of these factors. Instead, the Examiner alleged that Appellants' claim 26 does not "do an action statement" because Appellants' claim language of "identifying servers and placing the replicas of the data at the identified servers does not render a result." Final Office Action, page 2. Appellants submit that the Examiner's statement lacks merit.

Claim 26 recites "means for redistributing the one or more replicas to the identified one or more servers to at least one of increase reliability of the data, increase availability of the data, or increase bandwidth utilization in the system." The result clearly is that the replicas of the data are redistributed (i.e., stored) to the identified one or more servers. This result is clearly "useful" in the sense that storage of the replicas at the identified servers increases reliability of the data, increases availability of the data, and/or increases bandwidth utilization in the system. This result is also clearly "tangible" in the sense that particular replicas are stored at particular servers. See, e.g., Appellants' specification at paragraph 0073. This result is clearly "concrete" in the sense that data is redistributed to the identified one or more servers and this redistribution can be repeated. See, e.g., Appellants' specification at paragraphs 0072-0073. Accordingly, the Examiner's statement that claim 26 lacks a result is without merit.

Further, the Examiner indicated that Appellants "need to use language such as in dependent claim 4: use of 'identifying servers to store replicas of the data.'" Final Office Action, page 2. Appellants note that the language of "identifying servers to store replicas of the data" is similar to the language of "means for identifying one or more of the servers to which to move the one or more replicas" that is already recited in claim 26. Therefore, claim 26, under the Examiner's reasoning, should satisfy the requirements of 35 U.S.C. § 101. It appears that under

the guise of a rejection under 35 U.S.C. § 101, the Examiner is attempting to cure deficiencies in the Examiner's rejections under 35 U.S.C. § 102 and 103 by trying to force Appellants to amend the claims.

For at least these reasons, it is respectfully submitted that claim 26 recites statutory subject matter under 35 U.S.C. § 101. Reversal of the rejection of claim 26 is respectfully requested.

6. Claim 27.

Independent claim 27 is directed to a file system that comprises a plurality of servers configured to store replicas of chunks of data, and a master connected to the servers. The master is configured to select one or more of the replicas to redistribute based on utilization of the servers, identify one or more of the servers to which to move the selected one or more replicas, and move the selected one or more replicas to the identified one or more servers.

The Examiner rejected claim 27 under 35 U.S.C. § 101 as allegedly lacking patentable utility. Final Office Action, paragraph 5. The Examiner alleged that claim 27 merely identifies or defines "a data process in which to be manipulated without giving rise to a concrete, useful and tangible result." Final Office Action, page 3. Appellants traverse the rejection and submit that the Examiner's allegation falls short of establishing a prima facie case to deny patentability under 35 U.S.C. § 101.

The USPTO specifically states, in the "Interim Guidelines for Examination of Patent Applications for Patent Subject Matter Eligibility" (hereinafter "Guidelines"), the following:

To properly determine whether a claimed invention complies with the statutory invention requirements of 35 U.S.C. 101, USPTO personnel must first identify whether the claim falls within at least one of the four enumerated categories of patentable subject matter recited in section 101 (process, machine, manufacture or composition of matter).

* * *

The burden is on the USPTO to set forth a prima facie case of unpatentability. Therefore if the examiner determines that it is more likely than not that the claimed subject matter falls outside all of the statutory categories, the examiner must provide an explanation. ... If the examiner can establish a prima facie case that a claim does not fall into a statutory category, that does not preclude complete examination of the application for satisfaction of all other conditions of patentability. The examiner must further continue with the statutory subject matter analysis as set forth below.

Interim Guidelines For Examination of Patent Applications for Patent Subject Matter Eligibility, 1300 OG 142, § IV. B., November 22, 2005. Despite this clear obligation, the Examiner failed to make a determination of whether claim 27 falls within one of the four enumerated categories of patentable subject matter recited in Section 101. Appellants respectfully submit that claim 27 falls within one of the four enumerated categories of patentable subject matter recited in Section 101.

For example, claim 27 is directed to a file system that includes a plurality of servers and a master. Thus, claim 27 is directed to a "machine," which is one of the four enumerated categories of patentable subject matter recited in Section 101. The Examiner's allegation that claim 27 identifies or defines a data process lacks merit because claim 27 is directed to a file system and not a data process.

In light of the foregoing reasons, Appellants submit that the rejection of claim 27 under 35 U.S.C. § 101 is improper.

Furthermore, even assuming, arguendo, that claim 27 is not directed to one of the four enumerated categories of patentable subject matter recited in 35 U.S.C. § 101, the Guidelines further state:

While abstract ideas, natural phenomena, and laws of nature are not eligible for patenting, methods and products employing abstract ideas, natural phenomena, and laws of nature to perform a real-world function may well be. In evaluating whether a claim meets the

requirements of section 101, the claim must be considered as a whole to determine whether it is for a particular application of an abstract idea, natural phenomenon, or law of nature, rather than for the abstract idea, natural phenomenon, or law of nature itself.

Id. at § IV. C.

A proper analysis under 35 U.S.C. § 101 entails a determination of whether the claimed invention is a practical application of an abstract idea, law of nature, or natural phenomenon (judicial exceptions to 35 U.S.C. § 101). The Examiner has not even alleged that Appellants' claim 27 is directed to an abstract idea, natural phenomenon, or a law of nature. Therefore, the Examiner's allegation that claim 27 identifies or defines a data process falls short of establishing that the claim is directed to an abstract idea, a natural phenomenon, or a law of nature.

Nevertheless, claim 27 recites a practical application because it provides a transformation of an article to a different state or thing. The Guidelines state that if a claim provides a transformation or reduction of an article to a different state or thing, then the Examiner shall end the inquiry and find that the claim meets the statutory requirement of Section 101. Id. at IV.C.2.a. Claim 27 recites a master that, among other things, that moves (i.e., stores) the selected one or more replicas to the identified one or more servers, thus transforming a collection of data into replicas of data that are stored by one or more servers.

In light of this, Appellants submit that the rejection of claim 27 under 35 U.S.C. § 101 is again improper.

Even assuming, arguendo, that claim 27 does not provide a transformation, claim 27 provides a practical application that produces a useful, tangible, and concrete result. According to the Guidelines, in determining whether a claim provides a practical application that produces a useful, tangible, and concrete result, the Examiner should consider and weigh the following

factors: (1) the USPTO's official interpretation of the utility requirement provides that the utility of an invention has to be (i) specific, (ii) substantial and (iii) credible; (2) the tangible requirement does not necessarily mean that a claim must either be tied to a particular machine or apparatus or must operate to change articles or materials to a different state or thing; and (3) the process must have a result that can be substantially repeatable or the process must substantially produce the same result again. Id. at IV.C.2.b.

The Examiner did not consider any of these factors. Instead, the Examiner alleged that Appellants' claim 27 does not "do an action statement" because Appellants' claim language of "identifying servers and placing the replicas of the data at the identified servers does not render a result." Final Office Action, page 2. Appellants submit that the Examiner's statement lacks merit.

Claim 27 recites a master that is configured to "move the selected one or more replicas to the identified one or more servers." The result clearly is that the replicas of the data are moved (i.e., stored) at the identified one or more servers. This result is clearly "useful" in the sense that storage of the replicas at the identified servers maximizes data reliability and availability, and maximizes network utilization. See, e.g., Appellants' specification at paragraphs 0067 and 0073. This result is also clearly "tangible" in the sense that particular replicas are stored at particular servers. See, e.g., Appellants' specification at paragraph 0073. This result is clearly "concrete" in the sense that data is moved to the identified one or more servers and this movement can be repeated. See, e.g., Appellants' specification at paragraphs 0072-0073. Accordingly, the Examiner's statement that claim 27 lacks a result is without merit.

Further, the Examiner indicated that Appellants "need to use language such as in

dependent claim 4: use of 'identifying servers to store replicas of the data.'" Final Office Action, page 2. Appellants note that the language of "identifying servers to store replicas of the data" is similar to the language "identify one or more of the servers to which to move the selected one or more replicas" that is already recited in claim 27. Therefore, claim 27, under the Examiner's reasoning, should satisfy the requirements of 35 U.S.C. § 101. It appears that under the guise of a rejection under 35 U.S.C. § 101, the Examiner is attempting to cure deficiencies in the Examiner's rejections under 35 U.S.C. § 102 and 103 by trying to force Appellants to amend the claims.

For at least these reasons, it is respectfully submitted that claim 27 recites statutory subject matter under 35 U.S.C. § 101. Reversal of the rejection of claim 27 is respectfully requested.

B. The Rejection Under 35 U.S.C. § 102(b) Based on Jindal et al. (U.S. Patent No. 6,324,580) Should be Reversed.

The initial burden of establishing a prima facie basis to deny patentability to a claimed invention is always upon the Examiner. In re Oetiker, 977 F.2d 1443, 24 USPQ2d 1443 (Fed. Cir. 1992). For a proper rejection under 35 U.S.C. § 102, each and every element as set forth in the claim must be found in a single prior art reference. Verdegaal Bros. v. Union Oil Co. of California, 814 F.2d 628, 2 USPQ2d 1051 (Fed. Cir. 1987). Prior legal precedent requires that the identical invention be shown in as complete detail as is contained in the claim. Richardson v. Suzuki Motor Co., 868 F.2d 1226, 9 USPQ2d 1913 (Fed. Cir. 1989).

1. Claims 1 and 8.

Independent claim 1 is directed to a method for distributing data in a system that includes a plurality of servers. The method comprises identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers; and placing the replicas of the data at the identified servers.

Jindal et al. does not disclose or suggest the combination of features recited in claim 1. For example, Jindal et al. does not disclose or suggest identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers.

The Examiner alleged that Jindal et al. discloses these features and cited the Abstract, the Summary, column 4, lines 40-67, and column 5, lines 57-60, of Jindal et al. for support. Final Office Action, page 4. Appellants submit that these sections of Jindal et al. provide absolutely no support for the Examiner's allegation.

In the Abstract, Jindal et al. discloses a method for load balancing requests for a replicated service or application among a plurality of servers operating instances of the replicated service or application, and choosing a preferred server from the plurality of servers to receive a request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers, as required by claim 1. Instead, Jindal et al. clearly discloses something quite different in that Jindal et al. discloses choosing one of a plurality of servers that is operating an instance of a replicated service to process a request for the replicated service. Abstract.

In the Summary, Jindal et al. discloses a load balancing policy that specifies one or more factors to be used in determining a server that is to receive a client request. Column 2, lines 35-39. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers, as required by claim 1.

At column 4, lines 40-67, Jindal et al. discloses:

In a present embodiment of the invention, information concerning the operation of computer servers executing a replicated service is collected and processed to identify a preferred server (e.g., the server with the smallest load or shortest response time). Illustrative pieces of information that are collected include a server's response time, its distance from a central server (such as a name server providing DNS services), its operational status (e.g., whether it is up or down), etc.

For purposes of the present invention a replicated service is a service (e.g., web browsing, electronic mail) that is available on multiple servers. For example, an organization providing a service or application that is visited or invoked by numerous clients may employ several web servers to handle the requests. Each of the several servers is considered to operate a separate instance of the replicated service or application. Individual users may thus be routed to, and their requests satisfied by, any of the several servers.

The collected information is then analyzed and a preferred server is identified in accordance with a selected policy. In accordance with one illustrative policy, the preferred server is the server that is least-loaded. Another policy identifies the preferred server as being the closest server. After the preferred server is identified, subsequent requests for the replicated service or application are directed to that server. For example, in a web-browsing environment a DNS lookup table, or zone file, is updated to indicate that requests for the replicated service are to be routed to the preferred server.

In this section, Jindal et al. discloses that information is collected from the servers to identify the preferred server to receive a client request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers, as required by claim 1. Instead, Jindal et al.

discloses that a replicated service is simply a service that is available on multiple servers.

Column 4, lines 49-51. Jindal et al. does not disclose or suggest the manner in which the replicated service is stored on the servers.

At column 5, lines 57-63, Jindal et al. discloses:

In one embodiment of the present invention, the specific server identified in the zone file is determined according to a selected policy, as discussed below. Further, the server identified in zone file 104 is updated from time to time in accordance with the selected policy in order to distribute client requests among the servers offering the replicated service.

In this section, Jindal et al. discloses that a specific server, among servers offering a replicated service, is determined according to a selected policy. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers, as required by claim 1.

The Examiner also alleged that Jindal et al. discloses "the request for replicated service or application among a plurality of servers combine a central replicated monitor object . . . clearly illustrate the storing of replicas of the data as demonstrated in applicant's claimed language of identifying and storing the replicas of the data" and cited column 7, lines 34-47, of Jindal et al. for support. Final Office Action, page 2.

Appellants find the Examiner's allegation unclear. Jindal et al. discloses that a replicated service or application is associated with an individual monitor object that collects and saves status information, and a replicated monitor object that collects information from the individual monitor objects. Column 2, lines 55-67. Even assuming, for the sake of argument, that either the individual monitor object or the replicated monitor object can reasonably be equated to a

replica of data (a point that Appellants do not concede), nowhere does Jindal et al. disclose or remotely suggest identifying ones of the servers to store an individual monitor object or a replicated monitor object based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers, as would be required by claim 1.

At column 7, lines 34-48, Jindal et al. discloses:

Replicated monitor object (RMO) 204 retrieves the information collected by status objects from each IMO associated with one replicated service or application. Therefore, in the illustrated embodiment where each of servers 110, 112 and 114 operate a separate instance of a replicated service (e.g., web browsing), RMO 204 collects data from IMOs 202a, 202b and 202c. If the servers also offered another replicated service (e.g., electronic mail) or application, a second RMO would illustratively operate on nameserver 100 for the purpose of retrieving information concerning that service from a different set of IMOs. A replicated monitor object may also be known as a central monitor object due to its coordination role on behalf of a central server (e.g., nameserver 100) receiving multiple requests for a replicated service or application.

In this section, Jindal et al. discloses that the replicated monitor object retrieves information collected by the status objects from each individual monitor object associated with a replicated service or application. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers, as required by claim 1.

Jindal et al. also does not disclose or suggest placing the replicas of the data at the identified servers, as further recited in claim 1. The Examiner alleged that Jindal et al. discloses this feature and cited column 4, lines 40-67, of Jindal et al. for support. Final Office Action, page 4. Appellants submit that this section of Jindal et al. provides absolutely no support for the Examiner's allegation.

Column 4, lines 40-67, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses that information is collected from the servers to identify the preferred server to receive a client request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest placing the replicas of the data at the identified servers, as required by claim 1.

For at least these reasons, it is respectfully submitted that claims 1 and 8 are not anticipated by Jindal et al. under 35 U.S.C. § 102. Reversal of the rejection of claims 1 and 8 is respectfully requested.

2. Claims 2 and 3.

Dependent claim 2 recites identifying underutilized ones of the servers as candidates to store the replicas of the data. Jindal et al. does not disclose or suggest the combination of features recited in claim 2.

The Examiner alleged that Jindal et al. discloses the feature of claim 2 and cited the Abstract, the Summary, and column 4, lines 40-67, of Jindal et al. for support. Final Office Action, page 4. Appellants submit that these sections of Jindal et al. provide absolutely no support for the Examiner's allegation.

As noted above, the Abstract of Jindal et al. discloses a method for load balancing requests for a replicated service or application among a plurality of servers operating instances of the replicated service or application, and choosing a preferred server from the plurality of servers to receive a request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying underutilized ones of the servers as candidates to store the replicas of the data, as required by claim 2.

As noted above, the Summary of Jindal et al. discloses a load balancing policy that

specifies one or more factors to be used in determining a server that is to receive a client request.

Column 2, lines 35-39. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying underutilized ones of the servers as candidates to store the replicas of the data, as required by claim 2.

Column 4, lines 40-67, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses that information is collected from the servers to identify the preferred server to receive a client request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying underutilized ones of the servers as candidates to store the replicas of the data, as required by claim 2.

For at least these reasons, it is respectfully submitted that claims 2 and 3 are not anticipated by Jindal et al. under 35 U.S.C. § 102. Reversal of the rejection of claims 2 and 3 is respectfully requested.

3. Claim 4.

Dependent claim 4 recites identifying ones of the servers that have not been involved in a recent data distribution as candidates to store the replicas of the data. Jindal et al. does not disclose or suggest the combination of features recited in claim 4.

The Examiner alleged that Jindal et al. discloses the feature of claim 4 and cited the Abstract, the Summary, column 4, lines 40-67, and column 5, lines 57-60, of Jindal et al. for support. Final Office Action, page 4. Appellants submit that these sections of Jindal et al. provide absolutely no support for the Examiner's allegation.

As noted above, the Abstract of Jindal et al. discloses a method for load balancing requests for a replicated service or application among a plurality of servers operating instances of

the replicated service or application, and choosing a preferred server from the plurality of servers to receive a request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers that have not been involved in a recent data distribution as candidates to store the replicas of the data, as required by claim 4.

As noted above, the Summary of Jindal et al. discloses a load balancing policy that specifies one or more factors to be used in determining a server that is to receive a client request. Column 2, lines 35-39. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers that have not been involved in a recent data distribution as candidates to store the replicas of the data, as required by claim 4.

Column 4, lines 40-67, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses that information is collected from the servers to identify the preferred server to receive a client request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers that have not been involved in a recent data distribution as candidates to store the replicas of the data, as required by claim 4.

Column 5, lines 57-63, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses that a specific server is determined according to a selected policy. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers that have not been involved in a recent data distribution as candidates to store the replicas of the data, as required by claim 4.

For at least these reasons, it is respectfully submitted that claim 4 is not anticipated by Jindal et al. under 35 U.S.C. § 102. Reversal of the rejection of claim 4 is respectfully requested.

4. Claim 5.

Dependent claim 5 recites identifying system conditions that affect two or more of the servers, and identifying ones of the servers as candidates to store the replicas of the data based on the identified system conditions. Jindal et al. does not disclose or suggest the combination of features recited in claim 5.

The Examiner alleged that Jindal et al. discloses the features of claim 5 and cited column 4, lines 40-67, column 5, lines 57-60, and column 6, lines 32-46, of Jindal et al. for support. Final Office Action, page 4. Appellants submit that these sections of Jindal et al. provide absolutely no support for the Examiner's allegation.

Column 4, lines 40-67, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses that information is collected from the servers to identify the preferred server to receive a client request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying system conditions that affect two or more of the servers, or identifying ones of the servers as candidates to store the replicas of the data based on the identified system conditions, as required by claim 5.

Column 5, lines 57-63, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses that a specific server is determined according to a selected policy. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying system conditions that affect two or more of the servers, or identifying ones of the servers as candidates to store the replicas of the data based on the identified system conditions, as required by claim 5.

At column 6, lines 32-46, Jindal et al. discloses:

In this mode of operation, status objects 200a, 200b and 200c are invoked on nameserver 100 for the purpose of gathering information from servers 110, 112 and 114, respectively. The configuration and purpose of the status objects depend upon the policy that has been selected for choosing a preferred server. For example, where the selected policy requires

choosing the least-loaded server (e.g., that which has the fastest response time), each status object measures the response time of its associated server. Illustratively, this may be accomplished by issuing a Ping (or similar) command to the server and measuring the response time. As another example, where the selected policy requires choosing the closest server the status object is illustratively configured to measure the number of hops from nameserver 100 to the object's associated server.

In this section, Jindal et al. discloses choosing a least-loaded server or a closest server. Jindal et al. explicitly discloses that this server is selected to receive a client request. Column 2, lines 36-43. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying system conditions that affect two or more of the servers, or identifying ones of the servers as candidates to store the replicas of the data based on the identified system conditions, as required by claim 5.

For at least these reasons, it is respectfully submitted that claim 5 is not anticipated by Jindal et al. under 35 U.S.C. § 102. Reversal of the rejection of claim 5 is respectfully requested.

5. Claim 6.

Dependent claim 6 recites that a number of the replicas of the data stored by the servers is user-configurable. Jindal et al. does not disclose or suggest the combination of features recited in claim 6.

The Examiner alleged that Jindal et al. discloses the feature of claim 6 and cited the Abstract, the Summary, column 4, lines 40-67, and column 5, lines 57-60, of Jindal et al. for support. Final Office Action, page 5. Appellants submit that these sections of Jindal et al. provide absolutely no support for the Examiner's allegation.

As noted above, the Abstract of Jindal et al. discloses a method for load balancing requests for a replicated service or application among a plurality of servers operating instances of the replicated service or application, and choosing a preferred server from the plurality of servers

to receive a request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest that a number of the replicas of the data stored by the servers is user-configurable, as required by claim 6.

As noted above, the Summary of Jindal et al. discloses a load balancing policy that specifies one or more factors to be used in determining a server that is to receive a client request. Column 2, lines 35-39. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest that a number of the replicas of the data stored by the servers is user-configurable, as required by claim 6.

Column 4, lines 40-67, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses that information is collected from the servers to identify the preferred server to receive a client request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest that a number of the replicas of the data stored by the servers is user-configurable, as required by claim 6.

Column 5, lines 57-63, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses that a specific server is determined according to a selected policy. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest that a number of the replicas of the data stored by the servers is user-configurable, as required by claim 6.

For at least these reasons, it is respectfully submitted that claim 6 is not anticipated by Jindal et al. under 35 U.S.C. § 102. Reversal of the rejection of claim 6 is respectfully requested.

6. Claim 7.

Independent claim 7 recites a system for distributing chunks in a network that includes a plurality of servers. The system comprises means for selecting ones of the servers to store

replicas of the chunks based on at least one of utilization of the servers, prior chunk distribution involving the servers, or failure correlation properties associated with the servers; and means for storing the replicas of the chunks at the selected servers to at least one of increase reliability of the chunks, increase availability of the chunks, or increase bandwidth utilization in the system.

Jindal et al. does not disclose or suggest the combination of features recited in claim 7. For example, Jindal et al. does not disclose or suggest means for selecting ones of the servers to store replicas of the chunks based on at least one of utilization of the servers, prior chunk distribution involving the servers, or failure correlation properties associated with the servers.

The Examiner alleged that Jindal et al. discloses these features and cited the Abstract, the Summary, column 4, lines 40-67, and column 5, lines 57-60, of Jindal et al. for support. Final Office Action, page 5. Appellants submit that these sections of Jindal et al. provide absolutely no support for the Examiner's allegation for at least reasons similar to reasons given with regard to claim 1.

Jindal et al. also does not disclose or suggest means for storing the replicas of the chunks at the selected servers to at least one of increase reliability of the chunks, increase availability of the chunks, or increase bandwidth utilization in the system, as further recited in claim 7. The Examiner alleged that Jindal et al. discloses these features and cited the Abstract, the Summary, column 4, lines 40-67, column 5, lines 57-60, and column 6, lines 31-45, of Jindal et al. for support. Final Office Action, page 5. Appellants submit that these sections of Jindal et al. provide absolutely no support for the Examiner's allegation.

As noted above, the Abstract of Jindal et al. discloses a method for load balancing requests for a replicated service or application among a plurality of servers operating instances of

the replicated service or application, and choosing a preferred server from the plurality of servers to receive a request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest means for storing replicas of the chunks at the selected servers, let alone means for storing the replicas of the chunks to at least one of increase reliability of the chunks, increase availability of the chunks, or increase bandwidth utilization in the system, as required by claim 7.

As noted above, the Summary of Jindal et al. discloses a load balancing policy that specifies one or more factors to be used in determining a server that is to receive a client request. Column 2, lines 35-39. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest means for storing replicas of the chunks at the selected servers, let alone means for storing the replicas of the chunks to at least one of increase reliability of the chunks, increase availability of the chunks, or increase bandwidth utilization in the system, as required by claim 7.

Column 4, lines 40-67, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses that information is collected from the servers to identify the preferred server to receive a client request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest means for storing replicas of the chunks at the selected servers, let alone means for storing the replicas of the chunks to at least one of increase reliability of the chunks, increase availability of the chunks, or increase bandwidth utilization in the system, as required by claim 7.

Column 5, lines 57-63, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses that a specific server is determined according to a selected policy. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest means for storing replicas of the chunks at the selected servers, let alone means for storing the replicas of the chunks to at least one of increase reliability of the chunks, increase availability of the chunks, or increase

bandwidth utilization in the system, as required by claim 7.

Column 6, lines 32-46, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses choosing a least-loaded server or a closest server. Jindal et al. explicitly discloses that this server is selected to receive a client request. Column 2, lines 36-43. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest means for storing replicas of the chunks at the selected servers, let alone means for storing the replicas of the chunks to at least one of increase reliability of the chunks, increase availability of the chunks, or increase bandwidth utilization in the system, as required by claim 7.

For at least these reasons, it is respectfully submitted that claim 7 is not anticipated by Jindal et al. under 35 U.S.C. § 102. Reversal of the rejection of claim 7 is respectfully requested.

C. The Rejection Under 35 U.S.C. § 103(a) Based on Jindal et al. (U.S. Patent No. 6,324,580) in View of Narendran et al. (U.S. Patent No. 6,070,191) Should be Reversed.

The initial burden of establishing a prima facie basis to deny patentability to a claimed invention is always upon the Examiner. In re Oetiker, 977 F.2d 1443, 24 USPQ2d 1443 (Fed. Cir. 1992). In rejecting a claim under 35 U.S.C. § 103, the Examiner must provide a factual basis to support the conclusion of obviousness. In re Warner, 379 F.2d 1011, 154 USPQ 173 (CCPA 1967). Based upon the objective evidence of record, the Examiner is required to make the factual inquiries mandated by Graham v. John Deere Co., 86 S.Ct. 684, 383 U.S. 1, 148 USPQ 459 (1966). The Examiner is also required to explain how and why one having ordinary skill in the art would have been led to modify an applied reference and/or combine applied references to arrive at the claimed invention. Uniroyal, Inc. v. Rudkin-Wiley Corp., 837 F.2d

1044, 5 USPQ2d 1434 (Fed. Cir. 1988).

In establishing motivation, it has been consistently held that the requisite motivation to support the conclusion of obviousness is not an abstract concept, but must stem from the prior art as a whole to impel one having ordinary skill in the art to modify a reference or combine references with a reasonable expectation of successfully achieving some particular realistic objective. See, for example, Interconnect Planning Corp. v. Feil, 227 F.2d 1132, 227 USPQ 543 (Fed. Cir. 1985).

1. Claims 19 and 20.

Independent claim 19 is directed to a method for redistributing chunks of data in a system that includes a plurality of servers that store replicas of the chunks. The method comprises monitoring utilization of the servers; determining whether to redistribute any of the replicas; selecting one or more of the replicas to redistribute based on the utilization of the servers; selecting one or more of the servers to which to move the one or more replicas; and moving the one or more replicas to the selected one or more servers.

Neither Jindal et al. nor Narendran et al., whether taken alone or in any reasonable combination, discloses or suggests the combination of features recited in claim 19. For example, Jindal et al. and Narendran et al. do not disclose or suggest determining whether to redistribute any of the replicas.

The Examiner alleged that Jindal et al. discloses this feature and cited column 6, lines 31-45 and 56-64, of Jindal et al. for support. Final Office Action, page 6. Appellants submit that these sections of Jindal et al. provide absolutely no support for the Examiner's allegation.

Column 6, lines 32-46, of Jindal et al. is reproduced above. In this section, Jindal et al.

discloses choosing a least-loaded server or a closest server. Jindal et al. explicitly discloses that this server is selected to receive a client request. Column 2, lines 36-43. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest determining whether to redistribute any of the replicas, as required by claim 19.

At column 6, lines 56-64, Jindal et al. discloses:

Illustratively, for each replicated service (or application) that is to be monitored (i.e., that is subject to load balancing) on a server, a separate status object operates on nameserver 100. In addition, each status object illustratively performs a single function (e.g., determine response time, determine a server's distance from nameserver 100). In alternative embodiments of the invention, however, a single status object may monitor multiple servers or services and/or perform multiple functions.

In this section, Jindal et al. discloses a status object for each replicated service or application to be monitored. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest determining whether to redistribute any of the replicas, as required by claim 19.

The Examiner admitted that Jindal et al. does not disclose redistributing replicas. Final Office Action, page 6. This further supports Appellants arguments that Jindal et al. does not disclose or suggest determining whether to redistribute any of the replicas, as required by claim 19. In other words, if Jindal et al. does not disclose redistributing replicas, as admitted by the Examiner, then no reasonable argument can be made that Jindal et al. discloses determining whether to redistribute any of the replicas, as required by claim 19.

Nevertheless, the Examiner alleged that Narendran et al. discloses data distribution techniques for load-balanced fault-tolerant web access including redistributing replicated data from a failed server to achieve rebalance and cited column 12, lines 12-38, of Narendran et al. for support. Final Office Action, pages 6-7. Appellants submit that the disclosure of Narendran et al. does not support the Examiner's allegation and, in fact, *teaches away* from the features

recited in claim 19.

At column 12, lines 12-38, Narendran et al. discloses:

The load balance that is achieved through initial distribution of the documents may be disturbed during the operation of the server cluster due to the following changes: (1) a server failure; (2) changes in the access probabilities of the documents, i.e., some documents which were initially "hot" could become "cold" and vice-versa, and (3) changes in the capacity of the document servers. In the case of dedicated web servers, the first and second changes are generally more likely than the third. In the event any of these changes occur, it is desirable to be able to rebalance the server loads without any major reconfiguration of the system. More specifically, the rebalancing should be accomplished without redistributing documents among the servers, since this would generally involve overhead in moving the documents between servers and may affect the availability of the system. It is preferable to instead achieve rebalance by adjusting only the redirection probabilities used by the redirection server. As will be shown below, this rebalancing can be characterized as a network flow problem. It should be noted that the network flow approach is suitable for use in situations in which documents on a failed server are replicated and therefore available on another server or servers. In situations in which the documents on the failed server are not replicated, the above-described binning algorithm may be used to redistribute the documents from the failed server to achieve rebalance with a minimal amount of document movement.

In this section, Narendran et al. discloses rebalancing server loads by adjusting the redirection probabilities used by the redirection server when a server fails, changes in the access probabilities of the document occurs, or changes in the capacity of the document servers occurs. Nowhere in this section, or elsewhere, does Narendran et al. disclose or remotely suggest determining whether to redistribute any of the replicas, as required by claim 19. Instead, Narendran et al. discloses "the rebalancing should be accomplished without redistributing documents among the servers, since this would generally involve overhead in moving the documents between servers and may affect the availability of the system" (emphasis added). Column 12, lines 23-26. Accordingly, Narendran et al. specifically *teaches away* from redistributing documents (which the Examiner apparently equates to the replicas recited in claim 19).

Narendran et al. also discloses that in situations in which the documents on the failed server are not replicated, the algorithm could be used to redistribute the documents from the failed server to achieve rebalance with a minimal amount of document movement. Column 12, lines 33-38. In this situation, however, the documents cannot reasonably be interpreted as replicas since Narendran et al. specifically discloses that the documents are not replicated. Column 12, lines 33-35.

The Examiner further alleged that the "replicated system has to have the ability of redistribution. Jindal clearly discloses the replicated services include the redistribution based on utilization of server (i.e. 'referred server') as illustrated in the claimed language." Final Office Action, page 2. Appellants submit that the Examiner's allegation lacks merit. Nowhere does Jindal et al. disclose or remotely suggest redistribution, let alone determining whether to redistribute any of the replicas, as required by claim 19.

The Examiner also alleged that Narendran et al. discloses "redistribution or redirection of the data." Final Office Action, pages 2-3. Appellants submit that Narendran et al. does not disclose, and in fact teaches away from, determining whether to redistribute any of the replicas, as required by claim 19, for at least the reasons given above.

The Examiner further alleged that the "Examiner considers redistribution and redirection of data . . . as being movement of data from point A to point B and vice versa (as well as point C, D and F)." Final Office Action, page 3. Regardless of whether redistribution and redirection both involve moving data, the Examiner has not established that either Jindal et al. or Narendran et al. discloses or suggests determining whether to redistribute any of the replicas stored by the servers, as required by claim 19, for at least the reasons given above.

Jindal et al. and Narendran et al. also do not disclose or suggest selecting one or more replicas to redistribute based on the utilization of the servers, as further recited in claim 19. The Examiner alleged that Jindal et al. discloses this feature and cited column 6, lines 31-45 and 56-64, of Jindal et al. for support. Final Office Action, page 6. Appellants submit that these sections of Jindal et al. provide absolutely no support for the Examiner's allegation.

Column 6, lines 32-46 and 56-64, of Jindal et al. are reproduced above. Nowhere in these sections, or elsewhere, for reasons similar to reasons given above, does Jindal et al. disclose or remotely suggest selecting one or more replicas to redistribute, let alone selecting one or more replicas to redistribute based on the utilization of the servers, as required by claim 19. Narendran et al. also does not disclose or suggest this feature and actually teaches away from this feature for reasons similar to reasons given above.

Because Jindal et al. and Narendran et al., whether taken alone or in any reasonable combination, do not disclose or suggest determining whether to redistribute any of the replicas or selecting one or more of the replicas to redistribute based on the utilization of the servers, Jindal et al. and Narendran et al. cannot reasonably disclose or suggest selecting one or more of the servers to which to move the one or more replicas or moving the one or more replicas to the selected one or more servers, as further recited in claim 19.

For at least these reasons, it is respectfully submitted that claims 19 and 20 are patentable over Jindal et al. and Narendran et al., whether taken alone or in any reasonable combination, under 35 U.S.C. § 103. Reversal of the rejection of claims 19 and 20 is respectfully requested.

2. Claims 21 and 22.

Dependent claim 21 recites identifying underutilized ones of the servers as candidates to

which to move the one or more replicas. Jindal et al. and Narendran et al., whether taken alone or in any reasonable combination, do not disclose or suggest the combination of features recited in claim 21.

The Examiner alleged that Jindal et al. discloses the feature of claim 21 and cited column 6, lines 31-45 and 56-64, and column 8, line 53 - column 9, line 27, of Jindal et al. for support. Final Office Action, page 7. Appellants submit that these sections of Jindal et al. provide absolutely no support for the Examiner's allegation.

Column 6, lines 32-46, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses choosing a least-loaded server or a closest server. Jindal et al. explicitly discloses that this server is selected to receive a client request. Column 2, lines 36-43. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying underutilized ones of the servers as candidates to which to move the one or more replicas, as required by claim 21.

Column 6, lines 56-64, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses a status object for each replicated service or application to be monitored. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying underutilized ones of the servers as candidates to which to move the one or more replicas, as required by claim 21.

At column 8, line 53 - column 9, line 27, Jindal et al. discloses that a replicated service is offered on multiple servers, some of which are locally located and some of which are remotely located. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying underutilized ones of the servers as candidates to which to move the one or more

replicas, as required by claim 21.

For at least these reasons, it is respectfully submitted that claims 21 and 22 are patentable over Jindal et al. and Narendran et al., whether taken alone or in any reasonable combination, under 35 U.S.C. § 103. Reversal of the rejection of claims 21 and 22 is respectfully requested.

3. Claim 23.

Dependent claim 23 recites identifying ones of the servers that have not been involved in a recent redistribution as candidates to which to move the one or more replicas. Jindal et al. and Narendran et al., whether taken alone or in any reasonable combination, do not disclose or suggest the combination of features recited in claim 23.

The Examiner alleged that Jindal et al. discloses the feature of claim 23 and cited column 6, lines 31-45 and 56-64, and column 8, line 53 - column 9, line 27, of Jindal et al. for support. Final Office Action, page 7. Appellants submit that these sections of Jindal et al. provide absolutely no support for the Examiner's allegation.

Column 6, lines 32-46, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses choosing a least-loaded server or a closest server. Jindal et al. explicitly discloses that this server is selected to receive a client request. Column 2, lines 36-43. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers that have not been involved in a recent redistribution as candidates to which to move the one or more replicas, as required by claim 23.

Column 6, lines 56-64, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses a status object for each replicated service or application to be monitored. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the

servers that have not been involved in a recent redistribution as candidates to which to move the one or more replicas, as required by claim 23.

At column 8, line 53 - column 9, line 27, Jindal et al. discloses that a replicated service is offered on multiple servers, some of which are locally located and some of which are remotely located. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers that have not been involved in a recent redistribution as candidates to which to move the one or more replicas, as required by claim 23.

For at least these reasons, it is respectfully submitted that claim 23 is patentable over Jindal et al. and Narendran et al., whether taken alone or in any reasonable combination, under 35 U.S.C. § 103. Reversal of the rejection of claim 23 is respectfully requested.

4. Claim 24.

Dependent claim 24 recites determining failure correlation properties associated with the servers, and identifying ones of the servers based on the failure correlation properties as candidates to which to move the one or more replicas. Jindal et al. and Narendran et al., whether taken alone or in any reasonable combination, do not disclose or suggest the combination of features recited in claim 24.

The Examiner alleged that Narendran et al. discloses the features of claim 24 and cited column 6, lines 28-51, of Narendran et al. for support. Final Office Action, pages 7-8. Appellants submit that this section of Narendran et al. provides absolutely no support for the Examiner's allegation.

At column 6, lines 28-51, Narendran et al. discloses:

The server system 10 of FIG. 1 can be configured to provide graceful performance degradation in the presence of document server failures with minimum reconfiguration of

the system. For example, in embodiments in which documents are replicated and therefore available on more than one server, the redirection mechanism noted above can be used to configure the server system such that the initial distribution of the documents need not be changed and no document movement is required. When a given document server fails, a network flow based algorithm may be used to recompute the redirection probabilities to each of the replicas of each of the documents in order to approximately rebalance the load across the remaining document servers. It should be noted that the server cluster may have some of the documents replicated less than k times, because the illustrative system is not configured to generate and place new copies of documents that were on a failed server. Furthermore, the above-noted properties that only one replica will receive all requests for a document (in a heterogeneous cluster) or that replicas of a document have equal access rates (in a homogeneous cluster) will no longer hold true in the event of server failure. However, once the failed server is repaired, the initial redirection probabilities can be reinstated.

In this section, Narendran et al. discloses that when documents are replicated, the redirection mechanism can reconfigure the server system "such that the initial distribution of the documents need not be changed and no document movement is required" (emphasis added). In other words, Narendran et al. specifically *teaches away* from identifying ones of the servers based on the failure correlation properties as candidates to which to move the one or more replicas, as required by claim 24.

For at least these reasons, it is respectfully submitted that claim 24 is patentable over Jindal et al. and Narendran et al., whether taken alone or in any reasonable combination, under 35 U.S.C. § 103. Reversal of the rejection of claim 24 is respectfully requested.

5. Claim 25.

Dependent claim 25 recites deleting the one or more replicas from one or more of the servers, and instructing the selected one or more servers to copy the one or more replicas from another one or more of the servers. Jindal et al. and Narendran et al., whether taken alone or in any reasonable combination, do not disclose or suggest the combination of features recited in claim 25.

The Examiner alleged that Jindal et al. discloses the features of claim 25 and cited column 6, lines 31-45 and 56-64, and column 8, line 53 - column 9, line 27, of Jindal et al. for support. Final Office Action, page 8. Appellants submit that these sections of Jindal et al. provide absolutely no support for the Examiner's allegation.

Column 6, lines 32-46, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses choosing a least-loaded server or a closest server. Jindal et al. explicitly discloses that this server is selected to receive a client request. Column 2, lines 36-43. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest deleting the one or more replicas from one or more of the servers, or instructing the selected one or more servers to copy the one or more replicas from another one or more of the servers, as required by claim 25.

Column 6, lines 56-64, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses a status object for each replicated service or application to be monitored. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest deleting the one or more replicas from one or more of the servers, or instructing the selected one or more servers to copy the one or more replicas from another one or more of the servers, as required by claim 25.

At column 8, line 53 - column 9, line 27, Jindal et al. discloses that a replicated service is offered on multiple servers, some of which are locally located and some of which are remotely located. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest deleting the one or more replicas from one or more of the servers, or instructing the selected one or more servers to copy the one or more replicas from another one or more of the servers, as required by claim 25.

For at least these reasons, it is respectfully submitted that claim 25 is patentable over

Jindal et al. and Narendran et al., whether taken alone or in any reasonable combination, under 35 U.S.C. § 103. Reversal of the rejection of claim 25 is respectfully requested.

6. Claim 26.

Independent claim 26 is directed to a system for redistributing data in a network that includes a plurality of servers that store replicas of the data. The system comprises means for monitoring utilization of the servers; means for selecting one or more of the replicas to redistribute based on the utilization of the servers; means for identifying one or more of the servers to which to move the one or more replicas; and means for redistributing the one or more replicas to the identified one or more servers to at least one of increase reliability of the data, increase availability of the data, or increase bandwidth utilization in the system.

Neither Jindal et al. nor Narendran et al., whether taken alone or in any reasonable combination, discloses or suggests the combination of features recited in claim 26. For example, Jindal et al. and Narendran et al. do not disclose or suggest means for selecting one or more of the replicas to redistribute based on the utilization of the servers.

The Examiner alleged that Jindal et al. discloses this feature and cited column 6, lines 31-45 and 56-64, of Jindal et al. for support. Final Office Action, page 8. Appellants submit that these sections of Jindal et al. provide absolutely no support for the Examiner's allegation for at least reasons similar to reasons given with regard to claim 19.

Jindal et al. and Narendran et al. also do not disclose or suggest means for redistributing the one or more replicas to the identified one or more servers to at least one of increase reliability of the data, increase availability of the data, or increase bandwidth utilization in the system, as further recited in claim 26. The Examiner alleged that Jindal et al. discloses these features and

cited column 6, lines 31-45 and 56-64, and column 8, line 53 - column 9, line 27, of Jindal et al. for support. Final Office Action, page 8. Appellants submit that these sections of Jindal et al. provide absolutely no support for the Examiner's allegation.

Column 6, lines 32-46, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses choosing a least-loaded server or a closest server. Jindal et al. explicitly discloses that this server is selected to receive a client request. Column 2, lines 36-43. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest means for redistributing one or more replicas to the identified one or more servers, let alone means for redistributing the one or more replicas to at least one of increase reliability of the data, increase availability of the data, or increase bandwidth utilization in the system, as required by claim 26.

Column 6, lines 56-64, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses a status object for each replicated service or application to be monitored. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest means for redistributing one or more replicas to the identified one or more servers, let alone means for redistributing the one or more replicas to at least one of increase reliability of the data, increase availability of the data, or increase bandwidth utilization in the system, as required by claim 26.

At column 8, line 53 - column 9, line 27, Jindal et al. discloses that a replicated service is offered on multiple servers, some of which are locally located and some of which are remotely located. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest means for redistributing one or more replicas to the identified one or more servers, let alone means for redistributing the one or more replicas to at least one of increase reliability of the data, increase availability of the data, or increase bandwidth utilization in the system, as required by

claim 26.

Narendran et al. also does not disclose or suggest means for redistributing and actually *teaches away* from means for redistributing for reasons similar to reasons given above.

For at least these reasons, it is respectfully submitted that claim 26 is patentable over Jindal et al. and Narendran et al., whether taken alone or in any reasonable combination, under 35 U.S.C. § 103. Reversal of the rejection of claim 26 is respectfully requested.

7. Claim 27.

Independent claim 27 is directed to a file system that comprises a plurality of servers configured to store replicas of chunks of data, and a master connected to the servers. The master is configured to select one or more of the replicas to redistribute based on utilization of the servers, identify one or more of the servers to which to move the selected one or more replicas, and move the selected one or more replicas to the identified one or more servers.

Neither Jindal et al. nor Narendran et al., whether taken alone or in any reasonable combination, discloses or suggests the combination of features recited in claim 27. For example, Jindal et al. and Narendran et al. do not disclose or suggest a master that is configured to select one or more of the replicas to redistribute based on utilization of the servers.

The Examiner alleged that Jindal et al. discloses this feature and cited column 6, lines 31-45 and 56-64, of Jindal et al. for support. Final Office Action, page 9. Appellants submit that these sections of Jindal et al. provide absolutely no support for the Examiner's allegation for at least reasons similar to reasons given with regard to claim 19.

Jindal et al. and Narendran et al. also do not disclose or suggest a master that is configured to move the selected one or more replicas to the identified one or more servers, as

further recited in claim 27. The Examiner alleged that Jindal et al. discloses this feature and cited column 6, lines 31-45 and 56-64, and column 8, line 53 - column 9, line 27, of Jindal et al. for support. Final Office Action, page 9. Appellants submit that these sections of Jindal et al. provide absolutely no support for the Examiner's allegation for at least reasons similar to reasons given with regard to claim 19.

Narendran et al. also does not disclose or suggest a master to move one or more replicas and actually *teaches away* from a master to move one or more replicas for reasons similar to reasons given above.

For at least these reasons, it is respectfully submitted that claim 27 is patentable over Jindal et al. and Narendran et al., whether taken alone or in any reasonable combination, under 35 U.S.C. § 103. Reversal of the rejection of claim 27 is respectfully requested.

VIII. CONCLUSION

In view of the foregoing arguments, Appellants respectfully solicit the Honorable Board to reverse the Examiner's rejections of claims 1-8 and 19-27 under 35 U.S.C. §§ 101, 102, and 103.

To the extent necessary, a petition for an extension of time under 37 C.F.R. § 1.136 is hereby made. Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, to Deposit Account No. 50-1070 and please credit any excess fees to such deposit account.

Respectfully submitted,

HARRITY SNYDER, L.L.P.

/Paul A. Harrity/

Paul A. Harrity
Reg. No. 39,574

Date: June 1, 2007
11350 Random Hills Road
Suite 600
Fairfax, Virginia 22030
(571) 432-0800

IX. CLAIM APPENDIX

1. A method for distributing data in a system that includes a plurality of servers, the method comprising:

identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers; and

placing the replicas of the data at the identified servers.

2. The method of claim 1, wherein the identifying ones of the servers includes: identifying underutilized ones of the servers as candidates to store the replicas of the data.

3. The method of claim 2, wherein the underutilized servers are identified based on disk space usage below a determined amount.

4. The method of claim 1, wherein the identifying ones of the servers includes: identifying ones of the servers that have not been involved in a recent data distribution as candidates to store the replicas of the data.

5. The method of claim 1, wherein the identifying ones of the servers includes: identifying system conditions that affect two or more of the servers, and identifying ones of the servers as candidates to store the replicas of the data based on the

identified system conditions.

6. The method of claim 1, wherein a number of the replicas of the data stored by the servers is user-configurable.

7. A system for distributing chunks in a network that includes a plurality of servers, comprising:

means for selecting ones of the servers to store replicas of the chunks based on at least one of utilization of the servers, prior chunk distribution involving the servers, or failure correlation properties associated with the servers; and

means for storing the replicas of the chunks at the selected servers to at least one of increase reliability of the chunks, increase availability of the chunks, or increase bandwidth utilization in the system.

8. A file system, comprising:

a plurality of servers that store replicas of chunks; and

a master connected to the servers, the master being configured to:

identify one or more of the servers to store a replica of a chunk based on at least one of utilization of the servers, prior chunk distribution involving the servers, or failure correlation properties associated with the servers, and

place the replicas of the chunk at the identified one or more servers.

19. A method for redistributing chunks of data in a system that includes a plurality of servers that store replicas of the chunks, the method comprising:

monitoring utilization of the servers;
determining whether to redistribute any of the replicas;
selecting one or more of the replicas to redistribute based on the utilization of the servers;
selecting one or more of the servers to which to move the one or more replicas; and
moving the one or more replicas to the selected one or more servers.

20. The method of claim 19, wherein the utilization of the servers relates to an amount of free disk space available at the servers.

21. The method of claim 19, wherein the selecting one or more of the servers includes:

identifying underutilized ones of the servers as candidates to which to move the one or more replicas.

22. The method of claim 21, wherein the underutilized servers are identified based on disk space usage below a determined amount.

23. The method of claim 19, wherein the selecting one or more of the servers includes:

identifying ones of the servers that have not been involved in a recent redistribution as

candidates to which to move the one or more replicas.

24. The method of claim 19, wherein the selecting one or more of the servers includes:

determining failure correlation properties associated with the servers, and
identifying ones of the servers based on the failure correlation properties as candidates to which to move the one or more replicas.

25. The method of claim 19, wherein the moving the one or more replicas includes:
deleting the one or more replicas from one or more of the servers, and
instructing the selected one or more servers to copy the one or more replicas from another one or more of the servers.

26. A system for redistributing data in a network that includes a plurality of servers that store replicas of the data, the system comprising:
means for monitoring utilization of the servers;
means for selecting one or more of the replicas to redistribute based on the utilization of the servers;
means for identifying one or more of the servers to which to move the one or more replicas; and
means for redistributing the one or more replicas to the identified one or more servers to at least one of increase reliability of the data, increase availability of the data, or increase

bandwidth utilization in the system.

27. A file system, comprising:

a plurality of servers configured to store replicas of chunks of data; and

a master connected to the servers, the master being configured to:

select one or more of the replicas to redistribute based on utilization of the
servers,

identify one or more of the servers to which to move the selected one or more
replicas, and

move the selected one or more replicas to the identified one or more servers.

X. EVIDENCE APPENDIX

None

XI. RELATED PROCEEDINGS APPENDIX

None